

EUROPEAN PATENT APPLICATION

Application number: **88119577.0**

Int. Cl.4: **G06F 15/16**

Date of filing: **24.11.88**

Priority: **21.12.87 US 136055**

Date of publication of application:
28.06.89 Bulletin 89/26

Designated Contracting States:
AT BE CH DE ES FR GB GR IT LI NL SE

Applicant: **Honeywell Bull Inc.**
3800 W. 80th Street
Minneapolis Minnesota 55431(US)

Inventor: **Curley, John L.**
23, Willow Ridge Road
North Andover Massachusetts 01845(US)
Inventor: **Hirsch, Thomas S.**
24, Fox Run Road
Bedford Massachusetts 01730(US)
Inventor: **Penney, John C.**
3, Abbott Lane

Chelmsford Massachusetts 01824(US)
Inventor: **Reisch, Ileana S.**
53, Nathan Lane
Carlisle Massachusetts 01741(US)
Inventor: **Staplin, Theodore R., Jr.**
104, Garrison Road
Chelmsford Massachusetts 01824(US)
Inventor: **Wurz, David A.**
12, Bellevue Avenue
Winchester Massachusetts 01890(US)

Representative: **Altenburg, Udo, Dipl.-Phys. et al**
Patent- und Rechtsanwälte
Bardhele-Pagenberg-Dost-Altenburg
Frohwitter-Geissler & Partner Galileiplatz 1
Postfach 86 06 20
D-8000 München 86(DE)

Apparatus and method for a data processing system having a peer relationship among a plurality of central processing units.

A data processing system is disclosed in which a plurality of central processing units have access to all the system resources, i.e., have a peer relationship. During initialization of the data processing system, all the system resources are allocated to the individual central processing units according to a preselected distribution, the identification of available resources thereafter being stored in the files of the individual central processing units. During the operation of the data processing system, the resources can be reallocated by a predetermined procedure. The central processing units entering such a relationship are required to include apparatus and/or software procedures that prevent access to system resources not assigned thereto. A mail box procedure, using locations in the main memory unit permit communication between the central processing units and are used in the dynamic allocation of resources.




FIG. 4.

BACKGROUND OF THE INVENTION

1. Field of the Invention

This invention relates generally to data processing systems and, more particularly, to data processing systems having a plurality of central processing units.

2. Description of the Related Art

In order to increase the processing capability of data processing systems, one technique has been to couple additional central processing units to the system. The ability to select the number of central processing units in a data processing system permits an efficient matching of the capabilities of the system to the data processing requirements. Data processing units having a plurality of central processing units typically have one of two configurations. Referring now to Fig. 1A, a data processing system having a plurality of central processing units, according to a first implementation found in the prior art, is shown. The data processing system includes a plurality of central processing units 11-12 coupled to a system bus 19. The central processing units 11-12 perform the actual manipulation of data groups under control of operating and user software programs. The main memory unit 16, also coupled to the system bus 19, stores the data and program signal groups which are currently being used by the central processing units. The input/output units 14-15, coupled to the system bus 19, include devices for storage of large quantities of data and program signal groups, e.g., disk storage devices, terminals for the entry of data by system users, and communication devices for exchange of data and program groups with remote locations. The system bus 19 provides the principal path for the exchange of data and program groups between the components of the data processing system.

Referring next to Fig. 1B, a second implementation of a multiprocessor system, according to the related art, is shown. Generally, the same components are available to perform the processing functions as in Fig. 1A except that the components are coupled by a memory control unit 14 instead of by the system bus 19. The memory control unit 14 is typically an electronic switch providing the coupling of the data processing unit component in response to control signals. The memory control unit 14 can also provide functionality, such as conflict resolu-

tion, that would typically be distributed in the bus oriented data processing system.

The data processing systems of Fig. 1A and Fig. 1B are typically implemented in the related art such that the central processing units are homogeneous. In a homogeneous data processing system, the operating systems are the same or similar, the implementing apparatus is the same or similar and the operations performed on apparatus external to the data processing system is the same or similar. Even though the central processing units are homogeneous, substantial efforts are employed to prevent conflicts between the central processing units. For example, one of the central processing systems can be selected to allocate resources and tasks among the plurality of central processing units, thereby preventing conflicts for the resources by the plurality of programs that may be in current execution. The resources of the system are the storage devices, terminals, main memory locations and other data processing facilities to which a central processing unit has access for the purpose of performing the data processing functions. This relationship is generally referred to as the master/slave relationship because of the control asserted by the selected processor. However, some data processing systems can be designed wherein the central processing units, operating under control of the same operating system, can operate under as equal members (as contrasted with the master/slave relationship) of the data processing system. The following references provide examples of the way in which a plurality of central processing units can be incorporated in a data processing system without the master/slave relationships while still conflicts for system resources.

In U.S. Patent 3,631,405, issued December 28, 1971, entitled SHARING OF MICROPROGRAMS BETWEEN PROCESSORS and invented by G. S. Hoff and R. P. Kelly, two microprogrammed processing units share control elements that permit sharing of microprogram repertoires. By appropriate invocation of the operating system, the control signals from a first of the microprogrammed processing unit are transferred to the second microprogrammed processing unit. In fact, this configuration can best be described as a single processing unit with resources allocated by a supervisor controlled operating system. The use of a supervisor program as well as the coupling between the two processing units distinguishes this configuration from the peer processing unit relationship described in the present invention.

In U.S. Patent 4,131,941, issued December 26, 1978, entitled LINKED MICROPROGRAMMED PLURAL PROCESSOR UNIT and invented by H. L. Siegel, G. F. Muething, Jr., and E. J. Radkowski, a configuration of a plurality of processors is described that permits the processors to act independently or to be reconfigured so that a master/slave relationship can be invoked. The plurality of processors are linked together and, even when operating in a mode described as being independent, are not independent but subject to a supervisory control structure for configuration determination and for allocation of activity. Of course, the control of the allocation of activities implies the control of the allocation of resources. In addition, the data processing system described by this U.S. Patent, either has one operating system or a plurality of identical operating systems. The invention of the U.S. Patent appears to be best described as a single data processing system with a controllable configuration. The present invention is directed to data processing units that operate independently with different operating systems.

In U.S. Patent 4,200,930, issued on April 29, 1980, entitled ADAPTER CLUSTER MODULE FOR DATA COMMUNICATIONS SUBSYSTEM invented by R. L. Rawlings and R. D. Mathews, a host processing unit can have a plurality data communications subsystems coupled thereto for performing routine communications functions with incoming and outgoing signals. Although the data communications subsystems are capable, in case of a failure of the host processing unit, of continuing communications, the role of the host processing unit to the data communications subsystems is clearly that of a master/slave relationship. The peer processor relationship is not applicable because the data communications subsystems do not have access to all the resources available to the host processing unit.

In U.S. Patent Application Serial No. , filed , entitled MICROCOMPUTER SYSTEM WITH INDEPENDENT OPERATING SYSTEMS, invented by T. S. Hirsch, J. W. Stonier and T. O. Holtey, two processors, an LSI-6 processor with a MOD400 operating system and an Intel 8086 processor with either and MS-DOS or a CPM-86 operating share the processing responsibilities (a Motorola 6809 microprocessor is also included, but generally functions as an input/output controller). The LSI-6 processor has memory space that is not accessible to the Intel 8086 processor. In addition, the input/output operations performed by the 6809 microprocessor can be initiated only by the LSI-6 processor, so that the Intel 8086 has access to this resource only through the intervention of the LSI-6 processor, a form of the master/slave relationship.

In U.S. Patent Application Serial No. ,

entitled MULTIPROCESSOR SYSTEM ARCHITECTURE, invented by , the communication of two processors is described. In this application, the sharing of memory without interference is accomplished by controlling buses associated with each processor system. The buses are coupled to particular areas of memory and for one processor to access the memory dedicated to the second processor, the bus of the first processor is coupled to the bus of the second processor. Apparatus associated with each bus controls the ability of the other processor to access the bus, thereby effectively limiting access of each processor to the system resources.

More recently, interest has been demonstrated in data processing systems incorporating a plurality of central processing units, the central processing generally having non-homogeneous (generally incompatible) characteristics. The availability of non-homogeneous central processing units can be particularly advantageous to a system user providing the availability of a plurality of program repertoires. Ideally, all of the central processing units should have a peer relationship, i.e., should be capable of accessing all the data processing system resources without the benefit of protection auxiliary mechanisms described in relation to the related art and without having a master/slave relationship in which one central processing unit controls all the activity and allocation of resources. Many central processing systems do not have the necessary hardware and/or software functionality to enforce allocation of resources. Non-the-less, the peer relationship between central processing units is a desirable multiprocessor relationship, allowing easy expandability of the processing system.

A need has therefore been felt for technique that permits a plurality of (normally) incompatible central processing units to be coupled in a data processing system and to operate without conflict without requiring that a hierarchical relationship exist among the central processing units.

SUMMARY OF THE INVENTION

It is an object of the present invention to provide an improved data processing system that overcome the disadvantages of the prior art.

It is a feature of the present invention to provide an improved data processing system having a plurality of central processing units.

It is another feature of the present invention to provide an improved data processing system having a plurality of central processing units utilizing different operating systems.

It is a further feature of the present invention to

provide an improved data processing system in which a plurality of central processing units execute instructions under different operating systems, the central processing units having a peer relationship.

It is more particular feature of the present invention to provide apparatus permitting two central processing units of a data processing system to operate in a peer relationship while controlling the interaction between the central processing units.

It is still another feature of the present invention to permit a central processing unit to be coupled to a data processing system, in which all the central processing units have a peer relationship, even though the central processing unit being coupled uses an operating system and/or signal structures that are not compatible with the operating system 10 and/or signal structures of the data processing system.

The aforementioned and other features are accomplished, according to the present invention, by ensuring that a plurality of central processing unit/operating system combinations, to be coupled in data processing system, have appropriate mechanisms to prevent intentional or unintentional use of resources that have not been assigned to the central processing unit/operating system combination. With the internal mechanisms available to each central processing unit/operating system combination, each central processing unit/operating system combination can operate in a peer relationship, i.e., can have access to all the resources of the data processing system. During initialization procedures, the system resources are allocated to each of the central processing units. During operation of the data processing system, these resources can be dynamically reallocated among the central processing units.

These and other features of the invention will be understood upon reading of the following description along with the drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

Figure 1A is a block diagram of a bus oriented data processing system having a plurality of central processing units according to the related art, while Fig. 1B is a block diagram of a memory controller oriented data processing system having a plurality of central processing units according to the prior art.

Figure 2A illustrates the apparatus for implementing a peer processor relationship in a multiprocessor data processing system having operating system/central processing unit combinations

that are mutually incompatible, while Fig. 2B illustrates the partitioning of main memory storage for two operating system/central processing unit combinations having a peer processor relationship.

Figure 3 illustrates additional components that can be required for a multiprocessor data processing system having a peer relationship.

Figure 4 is a flow diagram illustrating how a user interacting with a data processing system by means of a first operating system can invoke processes requiring a second operating system.

DESCRIPTION OF THE PREFERRED EMBODIMENT

1. Detailed Description of the Figures

Referring to Fig. 2A, the principal components that permit two operating system/central processing units of a data processing system of multiprocessor data processing system that are generally not compatible to assume a peer relationship are illustrated. Each central processing unit has associated therewith apparatus, software procedures or a combination of apparatus and software programs that prevent the generation of addresses or attempts to access input/output devices are nominally unavailable to the central processing unit. This apparatus is illustrated as address generation apparatus 111 as part of central processing unit 11 and address generation security apparatus 121 as part of central processing unit 12. The main memory unit 15 is divided into a plurality of regions. Region 151 is a region reserved for the operating system controlling the operation of central processing unit 11, while region 152 is a portion of the main memory unit reserved for the operating system controlling operation of the central processing unit 12. The portion of the main memory unit 15 denoted by region 153 is reserved for the data/code for the central processing unit 11, while the region 154 is reserved for the data/code of central processing unit 12. Region 155 is a common region and accessible to both central processing unit and central processing unit 12.

Referring to Fig. 2A, more detail in the allocation of the storage of the main memory unit 15 is shown symbolically. The regions 151 and 152 reserved for Operating systems of the associated central processing unit each have two subregions (151A and 151B, and 152A and 152B, respectively) important for operation of a peer relationship. The subregions 151A and 152A provide a list of the resources, i.e., input/output devices and reserved main memory regions, reserved for the associated

central processing unit. Subregions 151B and 152B provide the procedure by which the two central processing systems can communicate. In the preferred embodiment, this communication is performed through subregion 155A of the portion of main memory unit 15 reserved for usage by both the central processing units 11 and 12 using a technique typically referred to as a data processing system mail box. In this technique, a message is left in the mail box 155A by a first central processing unit and a second central processing unit either reads the contents of the mail box periodically or is alerted by the first central processing unit that a message is available. Upon reading of the contents of the mail box 155A, the second central processing unit can make an appropriate response.

Referring next to Fig. 3, a portion of a data processing system, with two central processing units 11 and 12 is shown. Central processing unit 11 is directly coupled to the system bus 19. Central processing unit 12 is coupled to the system bus 19 by system bus interface unit 31. The system bus interface unit and the central processing unit 12 are coupled to dedicated memory unit 32.

Referring next to Fig. 4, the technique by which a first operating system procedure can be invoked by a second operating procedure is illustrated. In step 401, a second operating system process requires a process that can only be executed by the first processing system, i.e., an input/output operation. The second operating system interacts with the second operating system device driver module 402 as if the second operating system can implement the process. The second operating system driver module, instead of executing the process, applies the driver module to the processor interrupt mechanism driver module 405 by means of the second operating system processor interrupt mechanisms functions 403 and the shared main memory locations 404. Second operating system processor interrupt mechanism 403 and processor mechanism driver module 405 can exchange interrupt signals. The processor interrupt library routines 406 adapts the activity requested by the second operating system and applies the resulting process block to the first operating system input/output server 407. The input/output server 407 applies the appropriate signals to the first operating system communication and driver modules 408 to cause the process originally requested by the second operating system user to be executed.

2. Operation of the Preferred Embodiment

In a ideal multiprocessor configuration, all of the central processing units would have access to

all the resources of the data processing system without the requirement for auxiliary protection mechanisms insuring the partitioning of resources among the central processing units. Moreover, many operating system/central processing unit combinations have the procedures and/or the apparatus which can prevent attempts to access a resource assigned to another central processing unit, but may still be incompatible with a host data processing system. For example, the certain central processing unit activity, such as procedures for accessing selected input/output channels, may not be possible by the central processing unit to be coupled to the data processing unit. Similarly, the central processing unit to be coupled to the data processing system may exchange signal groups with the remainder of the data processing system and can use protocols and formats that are different from the coupled data processing system. Examples of these differences are widths of the system bus signal paths and whether system bus has individual groups of signal paths dedicated to different functions. Finally, conflicts can arise during initialization. For example, different central processing units can expect to find the requisite initialization procedures in overlapping region of the main memory. The present invention discloses the technique and the apparatus which can permit an additional operating system/central processing unit combination to be treated in a peer relationship with respect to the other operating system/central processing unit(s) of a data processing system.

In the preferred embodiment, the host central processing unit 11 is a Honeywell Bull MRX data processing system operating under control of the MOD 400 operating system. The MOD 400 operating system includes procedures that provide for the trustworthiness of the system. In order to provide the data processing system with the capability of executing the wide repertoire of user programs using the UNIX operating system, a central processing unit 12 was selected that operated under control of the UNIX operating system. However, for the implementation of the UNIX operating system available for executing user programs with the selected central processing unit, the requisite protection mechanisms were believed to be inadequate. The patent application entitled "APPARATUS AND METHOD FOR ALTERABLE RESOURCE PARTITIONING ENFORCEMENT IN A DATA PROCESSING SYSTEM HAVING CENTRAL PROCESSING UNITS USING DIFFERENT OPERATING SYSTEMS", cited above, provides apparatus and method for supplying the access protection mechanisms without modification of the data processing system or the central processing unit.

Although the peer processing relationship provides an equality between the central processing

units of the data processing system with respect to resources and activity, two situations must be provided for in order to accommodate special circumstances. The first of the special circumstances relates to the protocol used by the system to exchange data signals. Related to the system bus protocol is the organization of main memory. Typically, a "host" central processing unit 11 will be adapted to operate with a system bus and will not require additional manipulation of data groups when transferring the data groups to the system bus. However, the "guest" central processing unit, in general adapted to operate with different system bus protocol, will require a system bus interface unit 31 to permit the transfer of the signal groups between the central processing unit 12 and the system bus. The system bus interface unit 31, to the extent not possible by the central processing unit 12, can develop appropriate control signals and, where necessary, adjust the size of the data groups to be consistent with the remainder of the data processing system.

While the peer processing relationship can be utilized during normal operation, during an initialization of the data processing system as described above, one central processing unit must provide the control, such as setting up the resource tables for all the central processing units. In addition, for the configuration of the preferred embodiment described in the previous paragraph, the guest central processing unit during initialization is programmed to access main memory locations that overlap main memory locations accessed by the host central processing system and attempts to access memory locations that are not available to the main memory system. In order to avoid this potential problem, in the preferred embodiment of the present invention, a dedicated memory unit 32 is coupled to guest central processing unit 12. During initialization of the data processing system, initialization data for the guest central processing unit 12 is entered by a host central processing unit in the dedicated memory unit 32 at locations corresponding to the locations that are accessed by the guest central processing unit during initialization. Thereafter, when the guest central processing unit 12 is initialized, the memory locations in the dedicated memory 34 are addressed by the guest central processing unit and the requisite initialization data transferred to the guest central processing unit 12.

In order to accommodate differences in protocols and formats for signal groups between a host data processing system and a guest central processing unit, a system bus interface unit 31 can be coupled between the host data processing system and the guest central processing units. The system bus interface unit 31 includes the apparatus that converts data, instruction and control signal

groups from the data processing system (i.e., the system bus 19 in Fig. 3) into a format compatible with the signal format(s) with which the guest processing unit is designed to operate. The system bus interface unit 31 also has the apparatus to buffer the signal groups against differences in the system clock of the data processing unit and the system clock of the guest central processing unit. Similarly, the system bus interface unit has the apparatus for converting the data, instruction and control signal groups from the guest central processing unit into a format that can be used by and synchronized with the system bus. The conversion and synchronization of the signal groups transferred between the guest central processing unit and the host data processing system can be accomplished by techniques of the related art and are generally implemented specific.

With respect to procedures that are incompatible between the guest central processing unit and the host data processing system, the technique for executing the incompatible procedure is illustrated in Fig. 4. In essence, the execution of the procedure is performed by a central processing unit capable of the procedure execution under the direction of the central processing unit requiring the execution. By providing each central processing unit with the ability to execute incompatible procedures by executing the procedure by a second central processing unit, data processing system resources (e.g., input/output devices), typically inaccessible to certain central processing units, are available to all the central processing units of the data processing system.

Although the invention has been described with reference to two (incompatible) central processing units, it will be clear that the technique described can be used with a multiplicity of central processing units to provide a peer relationship among the central processing units.

The foregoing description is included to illustrate the operation of the preferred embodiment and is not meant to limit the scope of the invention. The scope of the invention is to be limited only by the following claims. From the foregoing description, many variations will be apparent to those skilled in the art that would yet be encompassed by the spirit and scope of the invention.

Claims

1. A data processing system comprising:
 - a plurality of resources;
 - a plurality of central processing units, each of said central processing units including mechanisms ensuring access to only selected resources, wherein at least two of said central processing units have

incompatible operating characteristics, each of said central processing units having files identifying said selected resources; and

compatibility means for permitting a central processing unit incompatible with a remainder of said data processing system to be coupled to said data processing system, said compatibility means permitting said incompatible central processing unit to access all of said system resources.

2. The data processing system of Claim 1 wherein said compatibility means includes an interface means for converting signal groups from said incompatible central processing unit to signal groups capable of being processed by said remainder of said data processing system and for converting signal groups from said remainder of said data processing unit to signal groups capable of being processed by said incompatible central processing unit.

3. The data processing unit of Claim 2 wherein said compatibility means includes procedure means for executing a procedure by a remainder of said data processing system in response to signals from said incompatible data processing system.

4. The data processing system of Claim 1 wherein said mechanisms ensuring address to selected resources includes means for verifying generated addresses are included in said selected resources.

5. The method of providing a data processing system having a plurality of central processing units, at least two of said central processing units being incompatible, wherein each of said plurality of data processing systems has a peer relationship with other of said plurality of central processing units, comprising the step of:

associating with each central processing unit a file identifying resources accessible by said associated central processing unit;

ensuring that every address generated by each data processing system is included in said file identifying accessible resources; and

permitting each central processing unit access to all resources of said data processing system.

6. The method of providing a peer relationship between central processing unit of a data processing unit of Claim 9 wherein said permitting step includes the step of:

for each incompatible central processing unit processing signal groups not capable of being processed by a remainder of said data processing system, converting signal groups provided by said remainder of said data processing system to said incompatible central processing unit to a format capable of being processed by said incompatible central processing unit and converting signal groups provided to said remainder of said data

processing system by said central processing unit to a data signal groups capable of being processed by a remainder of said data processing system.

7. A data processing system comprising:

5 data processing system resources including at least a main memory unit;

a first central processing unit/operating system combination capable of accessing said data processing system resources, said first combination including a first list of accessible resources and a first mechanism for ensuring an address generated by said first combination is included in said associated list of resources;

10 a second central processing unit/operating system combination incompatible with said first combination, said second combination including a second list of accessible resources and a first mechanism for ensuring that an address generated by said second combination is included in said second list;

15 interface means coupled between said second combination and said system resources for permitting said second combination to exchange signal with said system resources; and

20 allocation means for dynamically reallocating said data processing system resources between said first and said second file.

8. The data processing system of Claim 16 further comprising:

30 procedure means for permitting said second combination to execute an incompatible procedure by having said second combination instruct said first combination to execute said incompatible procedure.

9. The data processing system of Claim 17 wherein said main memory unit includes locations accessible to said first combination and to said second combination, wherein communication between said first combination and said second combination take place by means of said locations.

40 10. The data processing system of Claim 16 wherein said allocation means includes:

a first communication portion associated with said first combination;

45 a second communication portion associated with said second combination; and

locations in said main memory unit accessible to said first communication portion and said second communication portion.

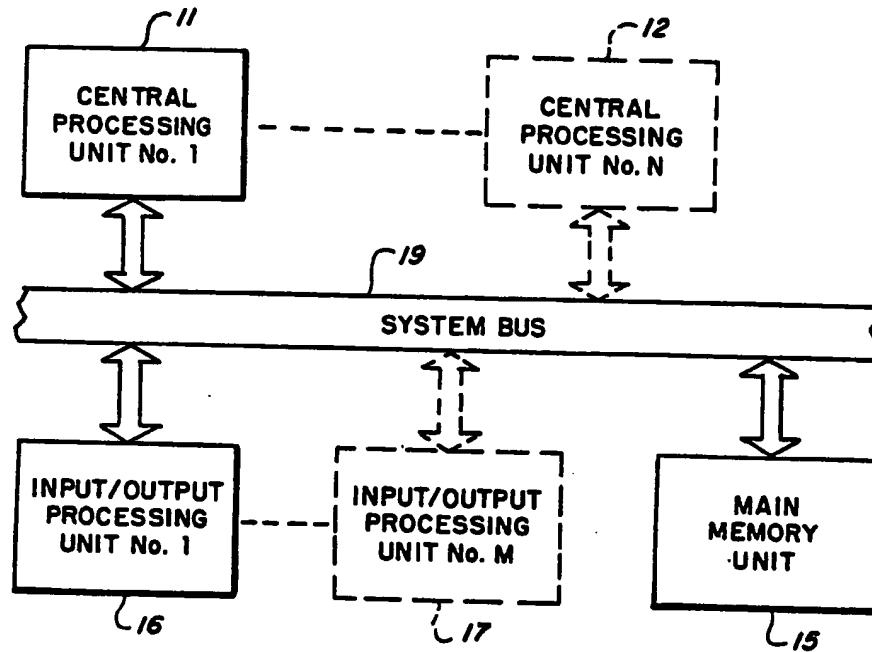


FIG. 1A.




FIG. 1B.




FIG. 2A.




FIG. 2B.




FIG. 3.

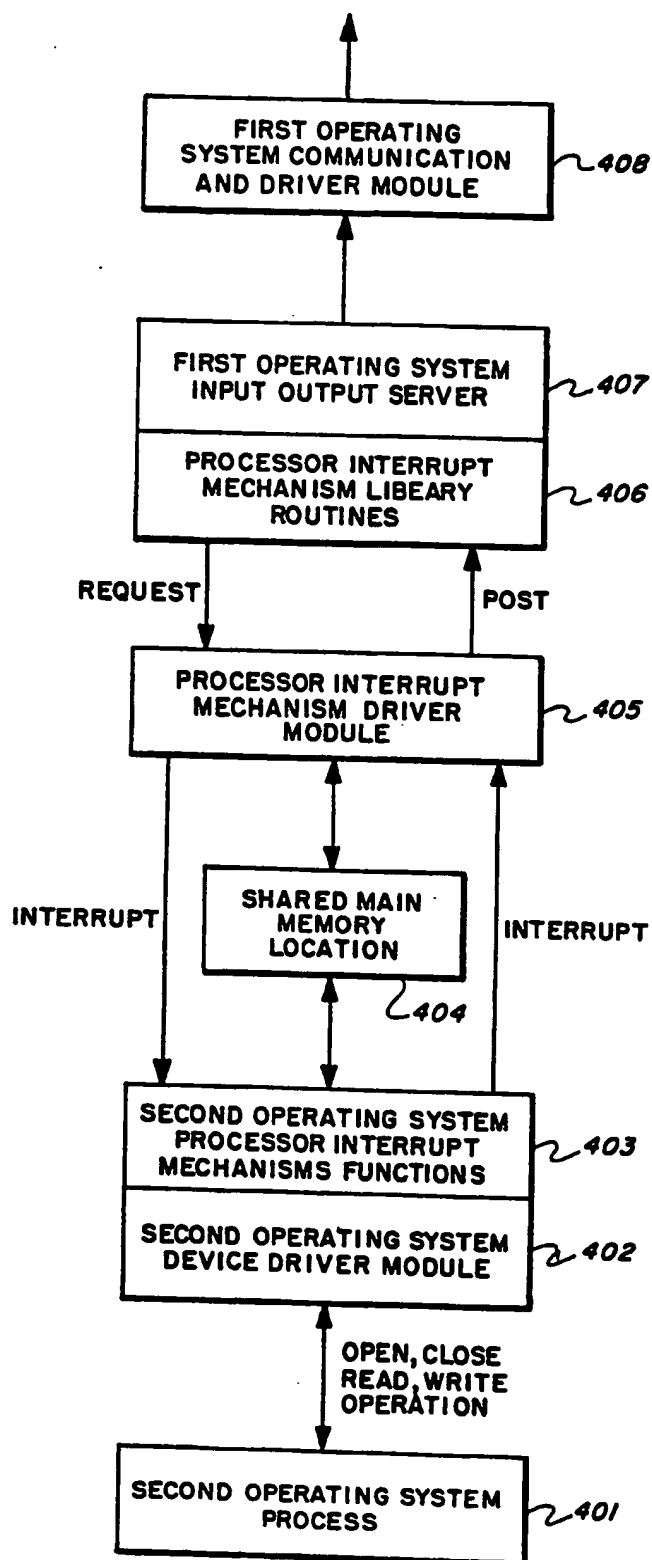


FIG. 4.